

AMENDMENTS TO THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application. Where claims have been amended and/or canceled, such amendments and/or cancellations are done without prejudice and/or waiver and/or disclaimer to the claimed and/or disclosed subject matter, and Applicant reserves the right to claim this subject matter and/or other disclosed subject matter in a continuing application or otherwise.

1. (previously presented) A filing system controlling block-level storage on a plurality of storage units, comprising:

a policy manager comprising at least one rule relating to block-level storage for a RAID level of protection for a file stored on the plurality of storage units, the RAID level of protection being selected from a plurality of RAID levels of protection, and at least one rule being based on an access pattern of files stored on the plurality of storage units, the filing system comprising information for each data block of the file indicating a number of other files in the filing system that require the data block for providing parity information for rebuilding each of the other files based on a parity calculation; and

an access manager providing the policy manager with information relating to access patterns of files stored on the plurality of storage units.

2. (original) The filing system according to claim 1, wherein the selected RAID level of protection is selected further based on size of the file.

3. (original) The filing system according to claim 1, wherein the selected RAID level of protection is selected further based on contents of the file.

4. (original) The filing system according to claim 1, wherein the selected RAID level of protection is selected further based on the name of the file and a location of the file in a name space of the filing system.

5. (original) The filing system according to claim 1, wherein at least two files are stored on the plurality of storage units having different RAID levels of protection.

6. (original) The filing system according to claim 1, wherein at least two files stored on a same storage unit have different RAID levels of protection.

7. (original) The filing system according to claim 1, wherein the information relating to access patterns of files is used for determining at least one RAID stripe size.

8. (currently amended) The filing system according to claim 1, wherein the information relating to access patterns of files is used for write coalescing data for storage on the plurality of storage units, and

wherein the filing system coalesces data in a partially full RAID stripe with data from another RAID stripe to make unused space available.

9. (original) The filing system according to claim 1, further comprising a RAID manager responsive to a rule contained in the policy manager by implementing the selected RAID level of protection for a file.

10. (original) The filing system according to claim 9, further comprising a RAID engine responding to the RAID manager by generating RAID redundancy-type information for the file.

11. (previously presented) The filing system according to claim 1, further comprising a space manager containing availability information for each storage block on the plurality of storage units.

12. (previously presented) The filing system according to claim 1, wherein at least one storage unit comprises a hard disk drive.

13. (previously presented) The filing system according to claim 1, wherein at least one storage unit comprises a random access memory device.

14. (previously presented) The filing system according to claim 1, wherein at least one storage unit comprises an optical drive.

15. (previously presented) A method of creating a file on a storage subsystem having a plurality of storage units, the method comprising:

receiving a request at a filing system to create a file on the plurality of storage units, the filing system comprising information for each data block of the file indicating a number of other files in the filing system that require the data block for providing parity information for rebuilding each of the other files based on a parity calculation;

querying a policy manager for at least one rule relating to block-level storage for a RAID level of protection for the file created on the plurality of storage units, the RAID level of protection being selected from a plurality of RAID levels of protection, and at least one rule contained in the policy manager being based on an access pattern of files stored on the plurality of storage units;

writing the file to the plurality of storage units based on the RAID level of protection selected for the file; and

maintaining metadata relating to a location of RAID information for the file within the filing system metadata information.

16. (original) The method according to claim 15, further comprising providing the policy manager with information relating to access patterns of files stored on the plurality of storage units.

17. (original) The method according to claim 15, wherein the selected RAID level of protection is selected further based on a size of the file.

18. (original) The method according to claim 15, wherein the selected RAID level of protection is selected further based on contents of the file.

19. (original) The method according to claim 15, further comprising storing at least two files on the plurality of storage units using different RAID levels of protection.

20. (original) The method according to claim 15, further comprising storing at least two files on a same storage unit using different RAID levels of protection.

21. (original) The method according to claim 15, further comprising determining at least one RAID stripe size based on the information relating to access patterns of files.

22. (currently amended) The method according to claim 15, further comprising write coalescing data for storage on the plurality of storage units based on the information relating to access patterns of files; and coalescing data in a partially full RAID stripe with data from another RAID stripe to make unused space available.

23. (original) The method according to claim 15, further comprising dynamically adjusting a RAID stripe size to match a filing system stripe and segment size based on at least one rule.

24. (original) The method according to claim 15, further comprising implementing the selected RAID level of protection for a file based on a rule contained in the policy manager.

25. (original) The method according to claim 24, further comprising generating RAID redundancy-type information for the file.

26. (original) The method according to claim 15, wherein at least one file stored on the plurality of storage units is stored as a store for filing system metadata information.

27. (original) The method according to claim 15, further comprising storing availability information for each storage block on the plurality of storage units.

28. (previously presented) The method according to claim 15, wherein at least one storage unit comprises a hard disk drive.

29. (previously presented) The method according to claim 15, wherein at least one storage unit comprises a random access memory device.

30. (currently amended) The method according to claim 15, wherein at least one storage unit ~~comprises~~comprises an optical drive.

31. (previously presented) A method of writing a file on a storage subsystem having a plurality of storage units, the method comprising:

determining at a filing system that a file stored on the plurality of storage units should be updated, the filing system comprising information for each data block of the file indicating a number of other files in the filing system that require the data block for providing parity information for rebuilding each of the other files based on a parity calculation;

querying a policy manager for at least one rule relating to block-level storage for a RAID level of protection for the file stored on the plurality of storage units, the RAID level of protection being selected from a plurality of RAID levels of protection, and at least one rule contained in the policy manager being based on an access pattern of files stored on the plurality of storage units;

writing the file to the plurality of storage units based on the RAID level of protection selected for the file; and

maintaining metadata relating to a location of RAID information for the file within the filing system metadata information.

32. (original) The method according to claim 31, wherein writing the file writes the file at the same place on the plurality of storage units that the file was located before the writing based on the selected RAID level of protection.

33. (original) The method according to claim 31, wherein writing the file writes the file at a different location on the plurality of storage units based on the selected RAID level of protection.

34. (currently amended) The method according to claim 31, further comprising providing the policy manager with information relating to access patterns of files stored on the plurality of storage units and coalescing data in a partially full RAID stripe with data from another RAID stripe to make unused space available.

35. (original) The method according to claim 31, wherein the selected RAID level of protection is selected further based on size of the file.

36. (original) The method according to claim 31, wherein the selected RAID level of protection is selected further based on contents of the file.

37. (original) The method according to claim 31, wherein the selected RAID level of protection is selected further based on a name of the file and a location of the file in a name space of the filing system.

38. (previously presented) The method according to claim 31, wherein at least one storage unit comprises a hard disk drive.

39. (previously presented) The method according to claim 31, wherein at least one storage unit comprises a random access memory device.

40. (previously presented) The method according to claim 31, wherein at least one storage unit comprises an optical drive.